Doursat, R. & Bienenstock, E. (2006b) Neocortical self-structuration as a basis for learning. *5th International Conference on Development and Learning (ICDL 2006)*, May 31-June 3, 2006, Indiana University, Bloomington, Indiana.

Neocortical Self-Structuration as a Basis for Learning

René Doursat
Department of Computer Science and Engineering
University of Nevada, Reno, Reno, NV 89557
doursat@unr.edu

Elie Bienenstock
Dept. of Neuroscience and Division of Applied Mathematics
Brown University, Providence RI 02912
elie@dam.brown.edu

Abstract – **Neocortical connectivity displays striking regularities that self-organize via epigenetic interactions with activity. We construe this self-structuration as a process of *spatiotemporal pattern formation* in a simple neural network model. Starting from random connections, ordered "synfire-chain" structures and wave-like correlations emerge simultaneously and reinforce each other through cooperation, while the global stability of the network is maintained by competition. Out of broad diffuse contacts and low stochastic firing, select synapses are strengthened and spike correlations increase. We suggest that the concurrent growth of multiple chains defines a *mesoscopic* scale of neural organization, and view the ontogenetic development of the nervous system as a "self-made tapestry" of synfire patterns, which could provide the basis for *compositionality* and learning.**

Index Terms – Synfire Chains, Self-Organization, Hebbian Learning, Spatiotemporal Patterns, Correlations, Compositionality.

I. INTRODUCTION

A. The Self-Structuration of Neocortex into Synfire Chains

The connectivity of neocortex displays striking regularities that account for its functional specialization. Synaptic contacts self-organize during pre- and post-natal development via epigenetic factors involving interactions between extrinsic and intrinsic sources of neural activity. Fundamentally, the growth of neocortex is guided by ongoing bioelectrical activity and competitive interactions. We propose here a neural network that models this self-structuration as a process of *spatiotemporal pattern formation*. We show the spontaneous and simultaneous emergence of ordered chains of synaptic connectivity together with a wave-like propagation of neuronal activity. This is based on two mechanisms: (a) Hebbian learning, one of the major principles underlying the development and tuning of the nervous system, and (b) the sensitivity of cortical neurons to nearly synchronous inputs amongst afferent connections. Synaptic contacts are rewarded by successful transmission events, as exemplified in spike-timing dependent plasticity (STDP). This work describes how a neural network may become spontaneously structured as a result of these principles, which are given a straightforward mathematical form.

The order emerging from our model assumes the shape of simple linear structures called "synfire chains" [1]. They consist of a sequence of synchronous groups $P_0 \rightarrow P_1 \rightarrow P_2 \rightarrow ...$, feed-forward connections and waves of activity (Fig. 1a). It is postulated that synfire chains could explain the preservation of accurately synchronized action potentials even in the presence of noise [2], as observed for example in the macaque MT area [3] and other recordings containing statistically significant *delayed* correlations in slices of mouse visual cortex [4]. These reproducible regimes of network dynamics can be construed in a generic way as "spatiotemporal patterns" of neural activity [5] or temporal modules dubbed "cortical songs" [6]. They correspond to specific constellations of spikes in raster views.

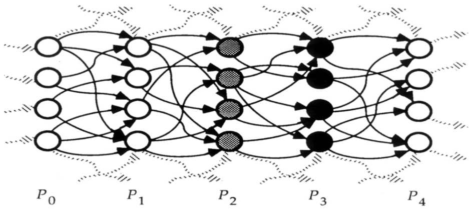

Fig. 1 Schema of a synfire chain. The geometry of the network has been unfolded along a temporal axis to make the linear structure appear clearly.

In our proposal, the main feature of synfire structuration resides in the fact that some synaptic connections are selected and reinforced to the detriment of others. An initially disordered network characterized by broad and diffuse connectivity and background stochastic activity undergoes a process of "focusing of innervation". Contacts that cooperate to sustain successful transmission events are rewarded, while activity evolves into a spatiotemporal regime characterized by high correlations among spike timings. The focusing expresses itself as the growth of a "long" but "narrow" structure. We shall see that, in a true self-organized fashion, connections and correlations reinforce each other through heterosynaptic *cooperation*, while the network's stability is maintained through a constraint of heterosynaptic *competition*.

B. Temporal Binding with Delays and Without Oscillations

This work also addresses the issue of an appropriate format of neural representation from a theoretical viewpoint. Our proposal is that the nervous system uses a *higher-order temporal code* $\langle x_i(t)\, x_j(t-\tau_{ij})\, x_k(t-\tau_{ik})\, ...\rangle$, instead of average firing rates $\langle x_i(t)\rangle$, to represent mental entities. The idea that accurate temporal relationships between neuronal firings may be used by the brain to express relationships has been formulated most notably by von der Malsburg [7]. Considerable interest in this approach has arisen in recent years, partly as a result of experimental data and theoretical studies suggesting synchronization of periodic firing as a mechanism for feature linking [8]. However, while using temporal synchrony for dynamical binding has become widely accepted [9], the oscillatory formulation is still debated. Moreover, phase locking is nearly always reported with phase zero. The representational power of the mechanism then appears to be rather limited. This study provides an alternative to the oscillatory-activity version of temporal coding. Specifically, we assume that entities to be bound (local features in a primary sensory cortex, lexical items in a language area, etc.) are represented by the activation of synfire chains, whose mutual binding is the counterpart to the phase locking of oscillators. One may say that the oscillator has been "unfolded" in space, so that a purely temporal pattern has now become a spatiotemporal one. Thus, we offer a simple explanation for the possible growth of synfire patterns and, second, suggest that these patterns may play an important role in the neural implementation of compositional cognitive functions. In our framework, accurate temporal relationships between the firing of neurons is used to express dynamical binding, based on synfire chains rather than oscillators.

C. *Mesoscopic Scale and Compositionality*

In this context, we postulate that synfire patterns might constitute the elementary components or "building blocks" at the microlevel of the mind's symbolic abilities, in particular the faculty of *compositionality* [10]—the assembling of elementary features into complex representations. Compositionality is fundamental to all cognitive functions. For example, the operative objects of perception and language can be described as symbols endowed with an internal structure, which allows them to assemble in multiple combinatorial ways and create a virtually infinite hierarchy of composite concepts. As discussed in the previous section, synfire structures precisely have the required properties to implement these constituents. Therefore, the theoretical claim carried by our model is that compositionality can arise from the gradual ontogenetic development of the nervous system during the early stages of synaptogenesis. In this, we join Chomsky's conception that the innate predisposition for the *faculty* of language (as opposed to any specific language) might involve a "growth" and maturation of the nervous system similar to a limb or an organ.

The present study was also motivated in response to classical criticisms stemming from the AI community against the lack of structured representations and symbolic support in connectionism. In neurally inspired systems, entities are embodied in the network's nodes and *activated* in an associationistic way. Conversely, logical systems define symbols that can be *composed* in a generative way, but do not possess a "microstructure" suited for perception or learning tasks. In our view, the missing link between these two levels is an intermediate or *mesoscopic* level of description, at a finer granularity than macroscopic symbols but higher structural complexity than microscopic neuronal activities. Here, neural groups such as synfire chains do not literally represent entities but rather provide the bricks for their construction. Metaphorically, it could be said that synfire structures are akin to "protein molecules", more complex than "atoms" (neurons) but more elementary than "cells" (cognitive representations). Our developmental model describes the creation of such a mesoscopic level by "crystallization" of the network into synfire modules. On this level, the modules interact and assemble to create wave synchronization via temporal binding (Fig. 8). Similarly to the vast majority of natural systems—physical, chemical, biological or social—we suggest that neocortex is a "self-made tapestry" pattern formation system [11], whose motifs combine dynamically to provide the basis for cognition and learning.

In the remainder of this article, Part II introduces a simple network model of the self-organized growth of synfire chains. Part III briefly analyzes the conditions of stability of the network in the random state and synfire-chain state, while Part IV describes in detail the iterative growth of a chain by accretion and consolidation of synchronous groups. Typical numerical results are presented in Part V, followed by a discussion of this work and future developments in Part VI.

II. A SIMPLE NETWORK MODEL

We consider a network of N excitatory neurons with binary values x_i representing spikes on the ms time scale. Synaptic weights w_{ij} vary by small increments on the same time scale as x_i. Time is discrete, in steps of roughly 1 ms, and connections have fixed transmission delays τ_{ij}. At each time t, the state of the network consists of action potentials $\mathbf{x}(t) = \{x_i(t)\}_{i=1\ldots N}$ and synaptic weights $\mathbf{w}(t) = \{w_{ij}(t)\}_{i,j=1\ldots N}$. This state evolves according to three laws: (a) neuronal activation, (b) synaptic plasticity and (c) intersynaptic competition.

Neurons obey a simple linear-nonlinear Poisson (LNP) dynamics, equivalent to the McCulloch & Pitts mean rate model transposed to the 1-ms timescale. The probability of activation of neuron j is given by $P[x_j(t) = 1] = \sigma_T(V_j(t) - \theta_j)$, where $V_j(t) = \sum_i w_{ij}(t)\, x_i(t - \tau_{ij})$ is the membrane potential of j at time t, θ_j its firing threshold and $\sigma_T(v) = 1 / (1 + \exp(-v/T))$, a sigmoidal step function. "Temperature" T controls the slope of the logistic function σ_T, i.e., the amount of noise in the system. The variation of connection weights depends on the fine temporal correlation between pre- and postsynaptic neurons. It is given by $w_{ij}(t) = w_{ij}(t-1) + b_{ij}(t)$, with $b_{ij}(t) = +\alpha$ for each $j \neq i$ such that $x_i(t - \tau_{ij}) = x_j(t) = 1$, and $b_{ij}(t) = -\beta$ if $x_i(t - \tau_{ij}) \neq x_j(t)$, where α and β are small positive numbers, typically of the order of .1 and .01, respectively. Thus, the effective rate of synaptic modification is much slower than that of the neuronal dynamics. The α-term is a schematic model of synaptic potentiation whereas the β-term represents synaptic depression. Presynaptic neurons must cooperate to increase the likelihood of successful transmission and receive synaptic reward. This fast synaptic plasticity is a form of Hebbian learning on the 1-ms time scale, and can also be viewed as a simplification of STDP, replacing the exponential curves with fixed increments.

The first two rules create a positive feedback in the network, whereby correlations and connections reinforce each other. To counterbalance this effect and prevent epilepsy, we introduce a third mechanism in the form of competition among synapses. We might want to impose that all outgoing ("efferent") and incoming ("afferent") weight sums be conserved at all times: $\sum_{j'} w_{ij'}(t) = \sum_{i'} w_{i'j}(t) = s_0$. Under such a constraint, the evolution of synaptic connections is better described as a *redistribution* rather than a creation of new contacts. For ease of calculation, we make this constraint a cost function $H(\mathbf{w}) = \gamma \sum_i [s_i^{\text{out}}(\mathbf{w}(t)) - s_0]^2 + \gamma \sum_j [s_j^{\text{in}}(\mathbf{w}(t)) - s_0]^2$, where $s_i^{\text{out}}(\mathbf{w}(t)) = \sum_{j'} w_{ij'}(t)$, $s_j^{\text{in}}(\mathbf{w}(t)) = \sum_{i'} w_{i'j}(t)$ and γ is of the order of .005. The synaptic rule becomes $w_{ij}(t) = w_{ij}(t-1) + b_{ij}(t) + c_{ij}(t)$, with $c_{ij}(t) = -(\partial H / \partial w_{ij})(\mathbf{w}(t-1) + \mathbf{b}(t))$. Finally, weights are clipped to stay inside [0, 1]. In summary, the network is driven by two forces: a positive feedback in the form of cooperation between activity and connectivity, and a corrective negative feedback in the form of competition among connections.

III. PRELIMINARY ANALYSIS OF ACTIVITY AND CONNECTIVITY

In this part and the next, we briefly analyze the behavior of the above model under simplified conditions, setting all delays τ_{ij} to a constant τ_0. Part V will also demonstrate synfire chain formation with nonuniform delays and inhibitory neurons in a diluted graph. Our first goal is tuning the network to a random activity mode with low average firing rate. Connectivity is broad and diffuse, with $w_{ij} \approx w_0 = s_0 / N$, for example, $N = 100$, $s_0 = 10$, $w_0 = .1$. Turning off the synaptic dynamics ($\alpha = \beta = \gamma = 0$), we look at the total activity in the network $n(t) = \sum_j x_j(t)$ and its expectation over several trials $\langle n(t)\rangle$. Denoting by n^* the average of $\langle n(t)\rangle$ over time, our goal is to obtain $0 < n^* << N$ to prepare conditions favorable to synfire growth. The neuronal activation dynamics yields $\langle n(t)\rangle = \sum_j P[x_j(t) = 1] = \sum_j \sigma_T(V_j(t) - \theta_j)$ and $\langle V_j(t)\rangle = w_0 \langle n(t - \tau_0)\rangle$. With $V_j(t) \approx \langle V_j(t)\rangle$ and $\theta_j \approx \theta_0$, we obtain mean field approximation $\langle n(t)\rangle \approx N\sigma_T(w_0\langle n(t - \tau_0)\rangle - \theta_0) \equiv \zeta(\langle n(t - \tau_0)\rangle)$, where ζ is an offset and rescaled version of σ_T, based on N, w_0 and θ_0. Thus, the average firing rate is a fixed point $n^* = \zeta(n^*)$. A quick graphical analysis reveals the existence of 1, 2 or 3 such points depending on the position of the sigmoid's center $\Omega = (\theta_0 / w_0, N / 2)$ and its slope at that point $\zeta'_\Omega = Nw_0 / 4T$ (Fig. 2). Placing Ω to the left of the diagonal, it means that ζ'_Ω must be steeper than the line from 0 to Ω to create two stable points, n_1^* near 0 and n_3^* near N. These conditions are $x_\Omega < y_\Omega \Leftrightarrow \theta_0 < s_0 / 2$ and $\zeta'_\Omega > x_\Omega / y_\Omega \Leftrightarrow T < \theta_0 / 2$. Compliant parameters as in Fig. 2 ensure a quick convergence to a stable low activity rate from a silent network $n(0) = 0$.

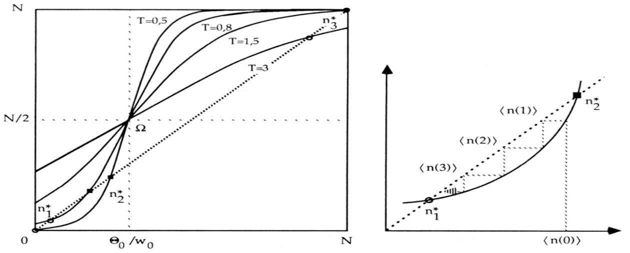

Fig. 2 Graphical study of mean activity in a random diffuse network. Parameters are $N = 100$, $w_0 = .1$ and $\theta_0 = 3$. Left: four sigmoids ζ with different values of T. For $T \leq .8$, there are 3 fixed points: n_1^* (near 0, stable), n_2^* (near θ_0 / w_0, unstable) and n_3^* (near N, stable). Right: we are interested in n_1^*, which is quickly reached from 0 ($n_1^* \approx 3.5$ for T = .8, or $n_1^* \approx .25$ for T = .5).

The same parameters allowing stable random activity in an unstructured network must also allow the stable propagation of a wave in a synfire chain. Before analyzing the creation of a synfire chain in the next sections, we examine how an existing chain can be maintained. Synfire group $P_{k\text{-}1}$ contacts P_k through strong efferent links w_+, while links from and to the rest of the network have near-zero value w_- (Fig. 1). Due to the sum rule, if the chain's width n_c (number of neurons in a group) is smaller than s_0, then $w_+ = 1$ and $w_- > 0$. Otherwise, if $n_c \geq s_0$, then $w_+ = s_0 / n_c \leq 1$ and $w_- = 0$. We shall see that only the latter case is stable, as pools keep growing until connections from and to the rest of the network are depleted. Let $m_k(t)$ be the number of active neurons in P_k at t and $\langle m_k(t)\rangle$ its expectation. Since P_k receives input only from $P_{k\text{-}1}$, $\langle m_k(t)\rangle \approx n_c\sigma_T(w_+\langle m_{k\text{-}1}(t - \tau_0)\rangle - \theta_0)$. Factoring in $w_+ = s_0 / n_c$ and $w_0 = s_0 / N$, we obtain the same recursive relation from group to group as in the random network: $\lambda\langle m_k(t)\rangle \approx \zeta(\lambda\langle m_{k\text{-}1}(t - \tau_0)\rangle)$, under a rescaling factor $\lambda \equiv N / n_c$. Thus, the same graphical analysis of Fig. 2 applies in a $n_c \times n_c$ square. However, for stable propagation to exist on the chain without dying out or exploding, two opposite conditions must be fulfilled: if $P_{k\text{-}1}$ is active at $t - \tau_0$, then P_k must be active at t, and if $P_{k\text{-}1}$ is inactive at $t - \tau_0$, then P_k must remain inactive at t. Thus, $m_k(t)$ must always stay close to $m_{k\text{-}1}(t - \tau_0)$, whether it is near n_c or near 0. This is why we need *both* stable fixed points n_1^* and n_3^*. The numerical values proposed in Fig. 2 satisfy both initial and final conditions. In this domain of parameters, the critical number of active P_0 neurons necessary and sufficient to reliably trigger a wave on the rest of the chain is given by n_2^*. This number is of course close to $\lambda x_\Omega = \theta_0 / w_+$.

Under these conditions, we switch on synaptic dynamics and calculate the effect of activity on weight variation. Since $\Delta\mathbf{w}(t)$ depends on $\mathbf{w}(t - 1)$ only via the efferent and afferent sums, which remain approximately constant, our calculation will be valid under any network structure, whether random or organized. We assume $\beta = 0$ to simplify the formulas. Denoting by P' and P the pools that fired at time $t - \tau_0$ and t, and by n' and n their size, the combination of Hebbian and sum rules yields $\Delta w_{ij}(t) = +\alpha(1 - 2\gamma(n + n'))$ for $i \in P'$ and $j \in P$; $-2\alpha\gamma n$ for $j \notin P$; $-2\alpha\gamma n'$ for $i \notin P'$; and 0 for all other connections. Therefore, to maintain a positive reward for the $P' \rightarrow P$ connections, the corrective competitive term $2\gamma(n + n')$ must remain small compared to 1. We adopt $\gamma \approx 1 / 2N$ in the remainder of this study. Fig. 3 shows $\Delta\mathbf{w}(t)$ as a bundle of reinforced links within the network and as a matrix. At the level of one connection, the competition term c_{ij} is of the order of n / N and small compared to the Hebbian term b_{ij}, of the order of 1. At the level of connection sums, $\sum_j c_{ij}$ and $\sum_j b_{ij}$ are both of the order of n, so $\mathbf{c}$ approximately counterbalances the effect of $\mathbf{b}$.

The variation $\Delta\mathbf{w}(t)$ represents the basic element of the network's self-structuration process. Alternatively denoting this matrix by $\mathbf{z}[P' \rightarrow P]$ and the set of active neurons $P(t)$, the synaptic dynamics reads $\mathbf{w}(t) = \mathbf{w}(t - 1) + \mathbf{z}[P(t - \tau_0) \rightarrow P(t)]$. Matrix $\mathbf{z}$ is a template that applies at every time step and under any degree of order in the network. Under uniform delays, $\mathbf{z}$ is nonzero only at instants multiple of τ_0. Thus, the global state of connection results from the iterated application of $\mathbf{z}$: $\mathbf{w}(k\tau_0) = \mathbf{w}(0) + \sum_{l = 0 \ldots k\text{-}1} \mathbf{z}[P(l\tau_0) \rightarrow P((l + 1)\tau_0)]$, showing that connectivity "deposits" or "builds up" in the wake of activity. Since, conversely, activity also tends to follow the path laid out by connectivity, the self-structuration of the network metaphorically resembles a "river" digging its own bed. The fate of $\mathbf{w}$ clearly depends on where the groups $P(t)$ pop up and two main cases can occur: (a) random activity perpetuating destruction, (b) ordered propagation fostering construction. If activity persists in the random regime, the **z** templates overlap inconsistently and cancel each other, therefore no structure can emerge. If on the contrary activity takes the form of a regular and reproducible wave, then **z** consistently accumulates between the same pools P_k and creates strong $P_k \to P_{k+1}$ connections: a chain can emerge and sustain itself by guiding the very wave that reinforces it. Both disorder and chain order are stable equilibria under the dual activity/connectivity dynamics.

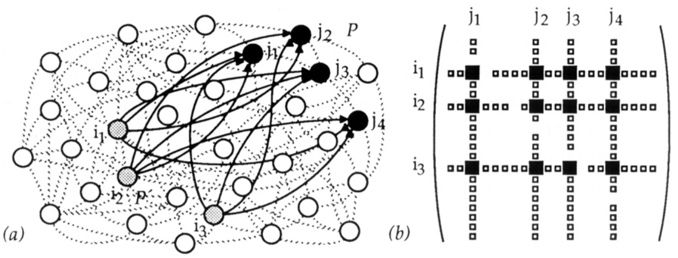

Fig. 3 Variation of network weights caused by activity. In this illustration, P' contains three neurons $i_1 \ldots i_3$, which were active at $t - \tau_0$ and P contains four neurons $j_1 \ldots j_3$, active at t. Connections from P' to P are reinforced while connections from P' to the outside, or from the outside to P, are weakened and the rest of the network does not change. Example of numerical values: with $N = 100$, $n' = 3$, $n = 4$, $\alpha = .1$, $\beta = 0$ and $\gamma = 1/2N = .005$, the four variations are +.093, –.004, –.003 and 0. (a) Network view of $\Delta \mathbf{w}(t)$. (b) Matrix view of $\Delta \mathbf{w}(t)$, in which reflexive contacts were prevented.

IV. PRINCIPLES OF NETWORK SELF-ORGANIZATION

After preparing the appropriate parametric conditions, we now examine the growth of a synfire chain. We also introduce a new singularity into the network—a special "seed" group of n_0 neurons, noted Q. These seed neurons tend to fire synchronously at a low average frequency, typically once every 10 or 20 τ_0, although not necessarily regularly. Through its repeated activation, Q will become the root of a new synfire chain.

A. *First iterations*

This section takes a closer look at the first steps of a chain's growth and shows how Q's pulses create a gradual differentiation in the landscape of activation probabilities. The chain will develop by aggregating new pools of synchronous neurons recursively. We set $N = 100$, $n_0 = 10$, $w_0 = .1$, $\theta_0 = 3$, $T = .5$, $\alpha = .1$, $s_0 = 10$. At first, potentials are near 0 and activation probabilities are uniformly low: $\sigma_T(-\theta_0) \approx n_1^* / N = .25\%$. When Q fires for the first time at t_0, it raises all activation probabilities to $\sigma_T(w_0 n_0 - \theta_0) = 1.8\%$. Consequently, an average of 2 neurons, noted P_1, fire at $t_0 + \tau_0$ and the weights immediately vary by $\mathbf{z}[Q \to P_1]$. Then, P_1 is typically too small to activate more neurons at $t_0 + 2\tau_0$, but the main effect is that $Q \to P_1$ weights have increased by about $\alpha = .1$ (neglecting the competition term), so their value is now $w_0 + \alpha = .2$. Therefore, when Q fires again at t_1, the landscape is not uniform anymore: neurons in P_1 now have a likelihood of activation $\sigma_T((w_0 + \alpha) n_0 - \theta_0) = 12\%$, whereas other neurons are still at $\sigma_T(w_0 n_0 - \theta_0) = 1.8\%$. The single increase of $+\alpha$ already reshaped the probability landscape because neurons in Q fired again in synchrony. However, it is still unlikely that the neurons in P_1 will fire again, despite their slight advantage. Rather, another 2 neurons from the rest of the network will be activated and P_1 now contains about 4 neurons. For a few more Q pulses, P_1 will continue growing and collecting neurons from the network, all characterized by a slightly increased likelihood of 12% of firing after the next Q pulse. Then, inevitably one or more neurons from P_1 will fire a *second* time after Q. This differentiates them further from the other neurons, as their connections from Q increase to $w_0 + 2\alpha = .3$ (in reality, a little less because of competition) and their next probability of firing jumps to $\sigma_T((w_0 + 2\alpha) n_0 - \theta_0) = 50\%$. The probability landscape is now partitioned into three groups: neurons in P_1 that fired after Q twice, once, or never, and respectively have 50%, 12% or 1.8% chance of firing again after Q.

B. *Creation of the First Synchronous Pool*

Neurons therefore differentiate on the basis of how many times they fired after Q. The fate of each neuron depends on its current degree of correlation with the seed group, i.e., the number of past $1 \to 1$ events. This is the core of the positive feedback loop: correlations increase connections, which encourage correlations. Once j is at 50%, it is very likely to be activated a 4th time by Q and reach $\sigma_T((w_0 + 3\alpha) n_0 - \theta_0) = 88\%$ soon. From this point on, j becomes a permanent member of P_1: $Q \to j$ connections reach the maximum $w_+ = s_0 / n_0$ and j is systematically triggered by Q, which maintains these connections. The $Q \to j$ connection pattern is extremely robust. Negligible in the beginning, competition becomes predominant in the later phase of P_1's formation. In the first phase, undifferentiated neurons become candidates to P_1's membership by firing once after Q. Then, in a second phase, these candidates compete to be recruited. Due to the sum rule, Q's efferent connections cannot afford to sustain all candidates, so neurons that fire again earlier "pull" synaptic contacts towards them, to the detriment of undecided neurons. Every time j is activated by Q, it earns a credit of connections, which it can gradually lose again if remains inactive too long, due to the gains made by other neurons. Each candidate is at a bifurcation point: a few timely post-Q activations propel its afferent weights upwards, whereas silence slowly takes them down. In sum, the aggregation of pool P_1 is "evolutionary": it first expands by *diversification* (driven by T) then shrinks and rounds up a final set of n_1 winners through *selection* (driven by s_0). To be able to durably recruit the first pool P_1, the seed group Q must contain a "critical mass" of at least θ_0 neurons. If $n_0 \leq s_0$, $Q \to j$ weights grow to 1, and $n_1 = s_0$. If $n_0 > s_0$, the weights reach $w_+ = s_0 / n_0$, and $n_1 = n_0$. Therefore, $n_1 = \max(n_0, s_0)$.

C. *Recursive Aggregation of a Synfire Chain*

Clearly, the same process from Q to P_1 repeats itself with P_1 playing the role of the seed and recruiting a new group P_2. The growth of synfire chain $Q \to P_1 \to P_2 \to \ldots$ is recursive and akin to the *growth by accretion of a crystal*. The latter also usually requires the presence of a singularity, acting as a seed. This accretion process is also not strictly iterative. The full period of P_1's development covers several of Q's firing events $t_0, t_1, \ldots, t_q$ (typically about 100), as neurons are recruited at different times and connections grow unequally on a first-come-first-served basis. Thus, P_1 does not wait for completion to start recruiting P_2: as soon as P_1 reaches critical mass θ_0, it is able to activate in turn new neurons downstream. The recruitment period of $Q \to P_1$ therefore greatly overlaps that of $P_1 \to P_2$, and so on. A growing chain typically lengthens before it widens, which gives it an aspect of a beveled head at the end of a mature trunk (Fig. 4). Several generations of immature groups along the growing tip of the chain compete to recruit neurons from the rest of the network. By recursive reasoning, P_{k-1} is always created before P_k, hence is larger and will be completed first. In the recruitment competition, older groups upstream maintain their advantage over younger groups downstream by instilling larger potentials into candidate neurons.

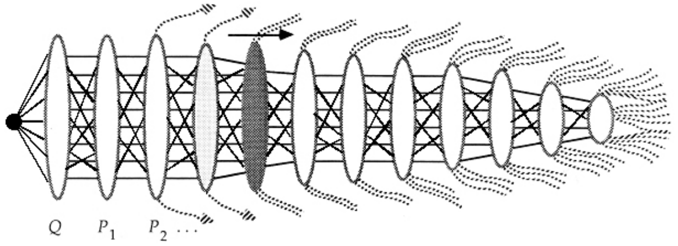

Fig. 4 Typical beveled profile of a growing synfire chain. The elongated ovals schematize synchronous pools in formation (height proportional to number of neurons) and the edges, strong $P_{k-1} \to P_k$ connections. A wave is currently propagating from P_3 to P_4. A mature "trunk" of constant width (here including P_2) ends in a "searching head" of decreasing width, due to largely overlapping periods of pool development. Immature efferent links (dashed lines) from unfinished pools compete to recruit new neurons from the network.

The activation of the seed group is an intrinsic inhomogeneity of the network. It is likely that at early stages of development some neurons send out stronger connections than others (black node in Fig. 4), providing adequate seeds for the process described here. This process is the intrinsic formation of connectivity structure under the influence of essentially endogenous bioelectrical activity. Finally, if several such seed groups coexist in the network and fire in an uncorrelated fashion, then several chains may simultaneously grow in the network (Fig. 8). These chains can interact, bind and support compositionality, as discussed in Part VI.

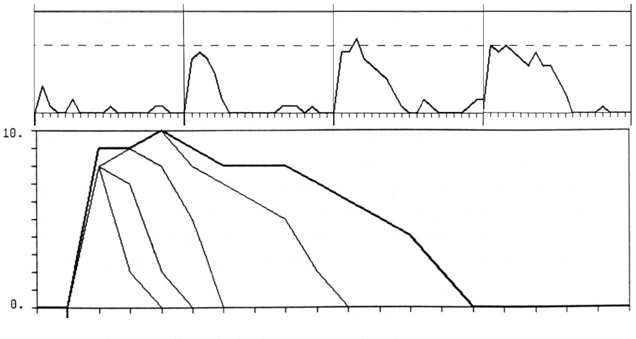

Fig. 5 A growing synfire chain in a network of 100 neurons. Parameters are $n_0 = 10$, $w_0 = .1$, $\theta_0 = 3$, $T = .5$, $\alpha = .1$, $s_0 = 10$. The firing period of the seed neuron is about 20. Top: four snapshots of the network's total activity following seed activations sampled at different times. Bottom: five other snapshots, superimposed in this view, revealing the growing profile of the chain (respectively at seed activations #10, 30, 50, 100 and 200).

V. NUMERICAL RESULTS

We conducted numerical simulations under the simplified assumptions of Parts III and IV, i.e., uniform delays and complete graph. Fig. 5 shows a full chain in development and Fig. 6 details the evolution of its connection values during the recruitment process (explanation in captions). The network also self-organizes under a more sophisticated model involving nonuniform transmission delays τ_{ij} and inhibition in a diluted graph (Fig. 7). The resulting structures are sequences of synchronous pools similar to synfire chains, elsewhere called "synfire braids" [5] or "polychronous groups" [12], in which longer delays connect nonconsecutive pools. A detailed inspection of the network (not shown) confirms that the strongest weights from group P_k to group P_{k+n} have indeed delay $n\tau_0$.

VI. DISCUSSION AND FUTURE WORK

We describe a process of self-organization characterized by the growth of a synfire chain from a seed group in a homogeneous network. Although we chose a specific set of equations and parameters for the neural activation and synaptic plasticity laws, it is readily seen that the principles underlying the model—the interplay of coincidence detection and synaptic reinforcement—are quite general, thus the described behav-

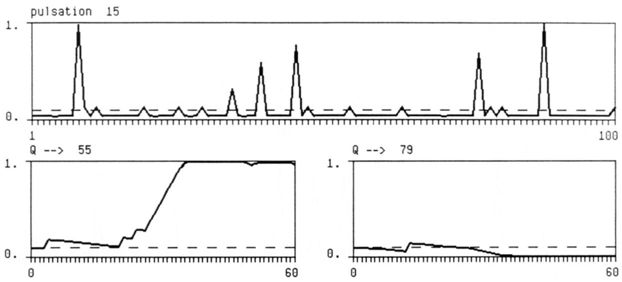

Fig. 6 Evolution of Q's outgoing connections during the recruitment process of the first pool P_1. Top (slice in space): snapshot of the landscape of weight values going from Q to the network, at the 15th seed activation (horizontal axis: neurons from 1 to 100). The neurons triggered by Q early on pull the connection weights towards them and diminish the chances of others to be recruited. From uniformly flat landscape at.1, 10 peaks appear, while the floor level drops to 0. Bottom (slice in time): fate of two specific sets of connections, $Q \to 55$ and $Q \to 79$ (horizontal axis: temporal sequence of seed activations). The discrete increments of $\alpha = .1$ caused by the neuron's activation are separated by periods of silence where the connections slowly decrease again because of competition.

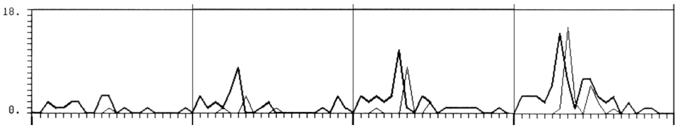

Fig. 7 Growth of a synfire braid in a network with delays and inhibition. There are 200 excitatory (E) and 40 inhibitory (I) neurons. Only the $E \to E$ synapses are plastic and their delays integer multiples of τ_0 in [1, 6]. $E \to I$ and $I \to E$ links form a fixed diluted graph (10%): they have constant values (≈ 1 and -5) and uniform delays τ_0. Four snapshots of the network's total activity following seed activations sampled at different times. Thick line: excitatory activity revealing an irregular synfire braid's profile. Thin line: inhibitory activity showing a delayed echo of the wave propagation. Inhibitory neurons are not a direct part of the chain but modulate its width in a quasi-oscillatory way: excitation increases inhibition, which in turn decreases excitation, etc.

ior should be largely independent of these choices. Neocortical structuration is a robust self-organization process, which is highlighted by the simplicity of our model—binary neurons, fixed weight increments and uniform graphs. Following the initial presentation of this work [13] (mentioned in [5]), the spontaneous formation of time-locked neuronal groups has also been demonstrated in a few other proposals based on more biologically detailed assumptions [14, 15, 12]. The latter in particular [12] is closely related to our work; using very similar principles but a rather different mathematical formulation, it provides an elegant demonstration of the self-organized growth of a large number of "polychronous groups" in an initially unstructured network. These spatiotemporal patterns are, essentially, short and thin synfire braids—one might say "braidlets." The braidlets described in [12] appear to be of a somewhat transient nature, compared to the longer, wider, and more persistent synfire structures that we observe. Synfire chain growth has also been reported in an "asynchronous regular" regime of activity characterized by intrinsic time shifts [16]; there, however, links develop mostly as a consequence of, and help stabilize, preexisting correlations.

Cortical connectivity exhibits a considerable amount of feedback. Via Hebbian plasticity, this creates a mutual reinforcement of excitatory activity and connectivity. Two different types of mechanisms may act to keep this situation under control. Our model incorporates one of them—*competition* between growing fibers converging onto, and emanating from a single neuron. This is a type of "reinforcement under constraints", which imposes a choice between all possible coherent patterns of connections and correlations. Here, this choice is largely random; it results from the stochastic or "thermodynamic" activation of neurons. Another mechanism is *inhibition*, which directly controls the global activity level in the network in an accurate way and on a faster time scale than synaptic competition. Much interest has recently been devoted to the study of "balanced networks", where the mean inhibitory input to a neuron cancels the mean excitatory input. In particular, it was suggested [17] that the propagation of a synchronous wave on the background of asynchronous activity can be achieved in a synfire chain containing inhibitory "shadow" pools to the excitatory ones, similarly to our Fig. 7.

At the core of perception and language is the capacity for handling composite entities never experienced before. Infinite productivity can arise, fundamentally, from combinatorics in a space of neural activity patterns. In the introduction, we outlined a possible use of synfire chains—or braids—as a substrate for compositional cognitive functions. This proposal [5] is based on the ability of such spatio-temporal patterns to *bind* with each other whenever synaptic interactions induce them to do so (Fig. 8). Synfire-chain binding is comparable to locking the phases of oscillators. We posit that weak excitatory synaptic interactions induce coupling between chains/braids, in the sense that it stabilizes the timing relationships across the chains. Such connections do not actually activate chains, but only ensure synchronization between already active chains. Compositionality in a system of synfire patterns has been recently examined in a model that included global inhibition [18]. It was proposed that inhibitory neurons help select among alternative wave couplings and prevent spurious synchronization. Another study [19] showed that synfire-type synchrony can also be controlled by balanced background synaptic input on each neuron: such input, if large enough, substantially increases the membrane conductance, hence makes synchronization more difficult to achieve.

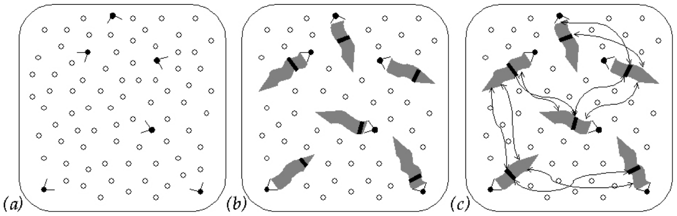

Fig. 8 Idealized sketch of a "self-made tapestry" of synfire chains. Multiple chains such as the one shown in Fig. 4 simultaneously develop from independent seed groups (a, b), then synchronize and bind (c). The wavy pointed rectangles represent growing synfire chains or braids in a rearranged view (not showing the fact that chains may also overlap and share neurons). The black stripes illustrate the current position of the waves of activation and the curvy arrows in (c) symbolize weak coupling connections between chains.

REFERENCES

[1] M. Abeles, *Local Cortical Circuits*, Berlin: Springer-Verlag, 1982.
[2] M. Diesmann, M.-O. Gewaltig, and A. Aertsen, "Stable propagation of synchronous spiking in cortical neural networks," *Nature*, vol. 402, pp. 529-533, 1999.
[3] P. Marsalek, C. Koch, and J. Maunsell, "On the relationship between synaptic input and spike output," *P. Natl. Acad. Sci. USA*, vol. 94, pp. 735–740, 1997.
[4] B.-Q. Mao, et al, "Dynamics of spontaneous activity in neocortical slices," *Neuron*, vol. 32, pp. 883-898, 2001.
[5] E. Bienenstock, "A model of neocortex," *Network*, vol. 6, 179-224, 1995.
[6] Y. Ikegaya, et al, "Synfire chains and cortical songs: temporal modules of cortical activity," *Science*, vol. 304, pp. 559-564, 2004.
[7] C. von der Malsburg, *The correlation theory of brain function*, Internal Report 81-2, Göttingen: Max Planck Institute for Biophys. Chem., 1981.
[8] G. Buzsáki, and A. Draguhn, "Neuronal oscillations in cortical networks," *Science*, vol. 304, pp. 1926-1929, 2004.
[9] A. L. Roskies, "The binding problem," *Neuron*, vol. 24, pp. 7-9, 1999.
[10] E. Bienenstock, "Composition," in *Brain Theory*, A. Aertsen and V. Braitenberg, Eds. Elsevier, pp. 269-300, 1996.
[11] P. Ball, *The Self-Made Tapestry: Pattern Formation in Nature*, New York: Oxford University Press, 1999.
[12] E. M. Izhikevich, "Polychronization: Computation with spikes," *Neural Computation*, vol. 18, pp. 245-282, 2006.
[13] R. Doursat, *Contribution à l'étude des représentations dans le système nerveux et dans les réseaux de neurones formels*, Ph.D. Thesis, Paris: Université Paris VI, 1991.
[14] J. Hertz, and A. Prügel-Bennett, "Learning short synfire chains by self-organization," *Network*, vol. 7, pp. 357-363, 1996.
[15] E. M. Izhikevich, J. A. Gally, and G. M. Edelman, "Spike-timing dynamics of neuronal groups," *Cereb. Cortex*, vol. 14, pp. 933-944, 2004.
[16] N. Levy, D. Horn, E. Meilijson, and E. Ruppin, "Distributed synchrony in a cell assembly of spiking neurons," *Neural Networks*, vol. 14, pp. 815-824, 2001.
[17] Y. Aviel, D. Horn, and M. Abeles, "Synfire waves in a small balanced network," *Neurocomputing*, vol. 58-60, pp. 123-127, 2004.
[18] M. Abeles, G. Hayon, and D. Lehmann, "Modeling compositionality by dynamic binding of synfire chains," *J. Comput. Neurosci.*, vol. 17, pp. 179-201, 2004.
[19] C. Trengrove, *Synfire Structures and Cognition: A Complex Systems Perspective*, Ph.D. Thesis, University of Technology Sydney, 2005.